arXiv:1509.08880v2 [stat.ML] 25 Nov 2015

Foundations of Coupled Nonlinear Dimensionality Reduction

Mehryar Mohri
Courant Institute and Google Research

Afshin Rostamizadeh
Google Research

Dmitry Storcheus
Google Research

Abstract

In this paper we introduce and analyze the learning scenario of *coupled nonlinear dimensionality reduction*, which combines two major steps of machine learning pipeline: projection onto a manifold and subsequent supervised learning. First, we present new generalization bounds for this scenario and, second, we introduce an algorithm that follows from these bounds. The generalization error bound is based on a careful analysis of the empirical Rademacher complexity of the relevant hypothesis set. In particular, we show an upper bound on the Rademacher complexity that is in $\widetilde{O}(\sqrt{\Lambda_{(r)}/m})$, where m is the sample size and $\Lambda_{(r)}$ the upper bound on the Ky-Fan r-norm of the associated kernel matrix. We give both upper and lower bound guarantees in terms of that Ky-Fan r-norm, which strongly justifies the definition of our hypothesis set. To the best of our knowledge, these are the first learning guarantees for the problem of coupled dimensionality reduction. Our analysis and learning guarantees further apply to several special cases, such as that of using a fixed kernel with supervised dimensionality reduction or that of unsupervised learning of a kernel for dimensionality reduction followed by a supervised learning algorithm. Based on theoretical analysis, we suggest a structural risk minimization algorithm consisting of the coupled fitting of a low dimensional manifold and a separation function on that manifold.

1 Introduction

Classic methods of linear dimensionality reduction assume that data approximately follows some low-dimensional linear subspace and aim at finding an optimal projection onto that subspace, i.e. Principle Component Analysis (PCA) [Pearson, 1901] and Random Projection [Hegde et al., 2008]. Nonlinear dimensionality reduction, also referred to as *manifold learning*, is a generalization of those linear techniques that aims at fitting a nonlinear low dimensional structure. Such manifold learning methods as Isometric Feature Mapping [Tenenbaum et al., 2000], Locally Linear Embedding [Roweis and Saul, 2000], and Laplacian Eigenmap [Belkin and Niyogi, 2001] are widely used as methods of nonlinear dimensionality reduction in machine learning, either to reduce the computational cost of working in higher-dimensional spaces, or to learn or approximate a manifold more favourable to subsequent learning tasks such as classification or regression. These algorithms seek to determine a nonlinear lower dimensional space by preserving various geometric properties of the input. However, it is not clear which of these properties would be more beneficial to the later discrimination stage. Since they are typically unsupervised techniques, they present a certain risk for the later classification or regression task: the lower-dimensional space found may not be the most helpful one for the second supervised learning stage and, in fact, in some cases could be harmful. How should we design manifold construction techniques to benefit most the subsequent supervised learning stage?

As shown by Figure 1, simply optimizing geometric properties may be detrimental the subsequent learning stage. To solve this problem, we consider an alternative scenario where the manifold construction step is not carried out *blindly*. We couple the task of nonlinear dimensionality reduction with the subsequent supervised learning stage. To do so, we make use of the known remarkable result that all of the manifold learning techniques already mentioned and many others are specific instances of the generic Kernel PCA (KPCA) algorithm for different choices of the kernel function [Ham et al., 2004]. More generally, all these methods can be thought of first mapping input vectors into a reproducing kernel Hilbert space and then conducting a low-rank projection within that space. Thus, our goal is to both learn a mapping as well as a projection taken from a parametric family as well as a hypothesis which is found in the low-dimensional space.

The main purpose of this paper is precisely to derive learn-

Copyright 2015 by the authors.

ing guarantees for this scenario, which we coin as *Coupled Nonlinear Dimensionality Reduction*, and to use those guaranteed as guidelines in the design of algorithms.

In practice, a user will often use a handful of different kernel functions and choose the one that is most effective according to measurements on a validation dataset. Instead, in this work, we argue that a more effective method is to allow a learning algorithm itself to choose a kernel function from a parametrized class. The idea of automatically selecting a kernel function has been explored in context of learning algorithms such as Support Vector Machines (SVM) [Lanckriet et al., 2004] and Kernel Ridge Regression (KRR) [Cortes et al., 2009] (see [Gönen and Alpaydın, 2011] and references therein for a more complete survey). To define the feature mapping, we will consider kernel families that consist of linear combinations of fixed *base kernel functions*. Such linear families have been analyzed extensively in the literature [Cortes et al., 2010, Kloft et al., 2011], however, mainly in the context of kernelized learning algorithms rather than dimensionality reduction techniques. Similarly, to define the projection, we will make use of the top-r eigenspace of a covariance operator that is defined as the linear combination of the covariances operators of the weighted base kernels. While some recent work has considered kernel learning in the setting of dimensionality reduction [Lin et al., 2011], to the best of our knowledge there has been no theoretical analysis or theoretical justification for the proposed algorithms. In this work, we provide the necessary theoretical analysis.

As mentioned above, within the setting of machine learning, dimensionality reduction is primarily used as a preprocessing step before regression/classification. For example, the recent work [Dhillon et al., 2013] illustrates the benefit of dimensionality reduction as preprocessing step by comparing the risk of OLS regression on reduced data to the risk of ridge regression on full data. They conclude that the risk of PCA-OLS is at most a constant factor of the risk of ridge regression, but can often be much less, as shown empirically.

Perhaps unsurprisingly, several empirical investigations have shown that tuning a dimensionality reduction algorithm in a coupled fashion, i.e. taking into account the learning algorithm that will use the reduced features, results in considerably better performance on the learning task [Fukumizu et al., 2004], [Gönen, 2014]. Despite this, the vast majority of existing theoretical analyses of dimensionality reduction techniques (even with fixed kernel functions) do not directly consider the learning algorithm that it will be used in conjunction with, and instead focus on the optimization of surrogate metrics such as maximizing the variance of the projected features [Zwald and Blanchard, 2005]. One exception is the work of [Mosci et al., 2007], which provides a generalization guarantee for hypothe-

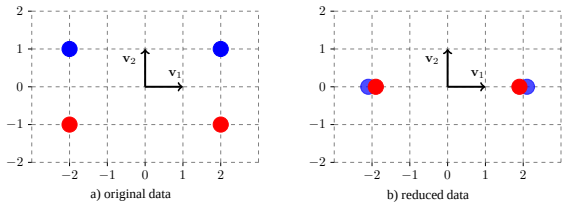

Figure 1: An example which illustrates that preserving only geometric properties may be detrimental to learning on reduced data. The original data in (a) has four points from blue and red classes. The eigenvectors of the covariance matrix are $\mathbf{v}_1 = \left(\begin{smallmatrix}1\\0\end{smallmatrix}\right)$ and $\mathbf{v}_2 = \left(\begin{smallmatrix}0\\1\end{smallmatrix}\right)$. Performing standard rank 1 PCA will project both blue and red points onto $\mathbf{v}_1$, thus merging them (as plotted in (b)). Any classification on the reduced data will necessarily incur a classification error of at least $1/2$.

ses generated by first conducting KPCA with a fixed kernel and then coupling with a regression model that minimizes squared loss. There is also a recent work of [Gottlieb et al., 2013], which derives generalization bounds based on Rademacher complexity for learning Lipschitz functions in a general metric space. They show that the intrinsic dimension of data significantly influences learning guarantees by bounding the corresponding Rademacher complexity in terms of dimension of underlying manifold and the distortion of training set relative to that manifold.

In our setting, we consider hypotheses which include both the KPCA dimensionality reduction step, with a *learned* linear combination kernel, as well as a linear model which uses the reduced features for a supervised learning task.

Although the hypothesis set we analyze is most naturally associated to a "Coupled Nonlinear Dimensionality Reduction" algorithm, which jointly selects both a kernel for nonlinear projection as well as a linear parameter vector for a supervised learning task, we note that this hypothesis set also encompasses algorithms that proceed in two stages, i.e. by first selecting a manifold and then learning a linear model on it.

The results of this paper are organized as follows: in the following section we outline the learning scenario, including the hypothesis class, regularization constraints as well as define notation. Section 3 contains our main result, which is an upper bound on the sample Rademacher complexity of the proposed hypothesis class that also implies an upper bound on the generalization ability of the hypothesis class. In Section 4 we show a lower bound on the sample Rademacher complexity as well as other quantities, which demonstrates a necessary dependence on several crucial quantities and helps to validate the design of the suggested hypothesis class. In Section 5 we provide a short discussion of the implications of our theoretical results, which leads us to Section 6, where develop an algorithm for the

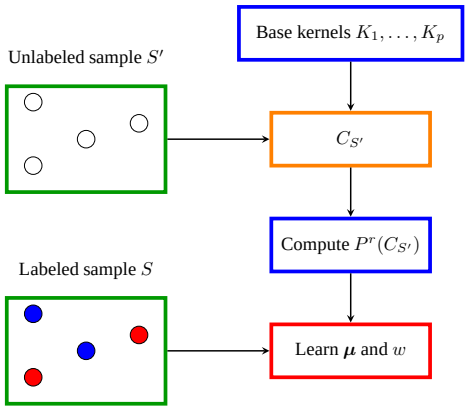

Figure 2: Flow chart illustrating the Coupled Dimensionality Reduction learning scenario.

coupled fitting of a kernel and a separation function.

2 Learning scenario

Here, we describe the learning scenario of supervised dimensionality reduction. Let $\mathcal{X}$ denote the input space. We assume that the learner receives a labeled sample of size m, $S = ((x_1, y_1), \ldots, (x_m, y_m))$, drawn i.i.d. according to some distribution $\mathcal{D}$ over $\mathcal{X} \times \{-1, +1\}$, as well as an unlabeled sample $U = (x'_1, \ldots, x'_u)$ of size u, typically with $u \gg m$, drawn i.i.d. according to the marginal distribution $\mathcal{D}_{\mathcal{X}}$ over $\mathcal{X}$.

We assume that the learner has access to p positive-definite symmetric (PDS) kernels $K_1, \ldots, K_p$. Instead of requiring the learner to commit to a specific kernel K defining KPCA with a solution subsequently used by a classification algorithm, we consider a case where the learner can define a dimensionality reduction solution defined based on $\sum_{k=1} \mu_k K_k$, where the non-negative mixture weights μ_k are chosen to minimize the error of the classifier using the result of the dimensionality reduction.

Given p positive-definite symmetric (PDS) kernels $K_1, \ldots, K_p$ and a vector $\boldsymbol{\mu} \in \mathbb{R}^p$ with non-negative coordinates, consider a set of weighted kernel functions $\{\mu_1 K_1, ..., \mu_p K_p\}$, where each $\mu_k K_k$ has its reproducing space $\mathbb{H}_{\mu_k K_k}$. The unlabeled sample U is used to define the empirical covariance operator of each weighted base kernel $\mu_k K_k$, denoted as $C_{U,\mu_k K_k} \colon \mathbb{H}_{\mu_k K_k} \to \mathbb{H}_{\mu_k K_k}$. Given a set of covariance operators $\{C_{U,\mu_k K_k}\}_{k=1}^p$ we define an operator $C_{U,\boldsymbol{\mu}} = \sum_{k=1}^p C_{U,\mu_k K_k}$ that acts on the sum of reproducing spaces $\mathbb{H}_{\boldsymbol{\mu}} = \mathbb{H}_{\mu_1 K_1} + ... + \mathbb{H}_{\mu_p K_p}$.

Let $P_r(C_{U,\boldsymbol{\mu}})$ denote the rank-r projection over the eigenspace of $C_{U,\boldsymbol{\mu}}$ that corresponds to the top-r eigenvalues of $C_{U,\boldsymbol{\mu}}$ denoted as $\lambda_1(C_{U,\boldsymbol{\mu}}) \geq \ldots \geq \lambda_r(C_{U,\boldsymbol{\mu}})$.[1] Let

[1] Although it is also possible provide results in an even more general setting, we make the assumption that the chosen r satisfies $\lambda_r(C_{U,\boldsymbol{\mu}}) \neq \lambda_{r+1}(C_{U,\boldsymbol{\mu}})$ in order to simplify the presentation. Note, assumption this is always satisfied if the eigenvalues are simple.

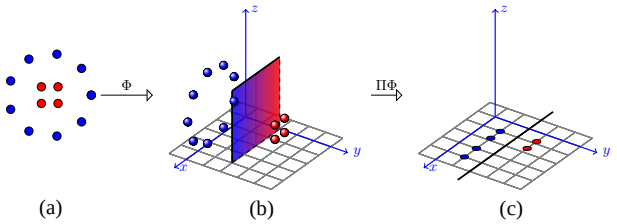

Figure 3: Illustration of the coupled learning problem. (a) Raw input points. (b) Points mapped to a higher-dimensional space where linear separation is possible but where not all dimensions are relevant. (c) Projection over a lower-dimensional space preserving linear separability.

$\Phi_{\mu_k K_k} \colon \mathcal{X} \to \mathbb{H}_{\mu_k K_k}$ denote the feature mapping associated to $\mu_k K_k$, specifically for each $x \in \mathcal{X}$, we have the function $\Phi_{\mu_k K_k}(x) = \sqrt{\mu_k} K_k(x, \cdot)$. Define $\Phi_{\boldsymbol{\mu}} \colon \mathcal{X} \to \mathbb{H}_{\boldsymbol{\mu}}$ as $\Phi_{\boldsymbol{\mu}} = \sum_{k=1}^p \Phi_{\mu_k K_k}$. This parametrized projection is used to define rank-r feature vectors (functions) $P_r(C_{U,\boldsymbol{\mu}})\Phi_{\boldsymbol{\mu}}(x)$.

From this point onward, in to avoid intricate notation, we will not not explicitly indicate the dependence of $\mathbb{H}_{\boldsymbol{\mu}}$, $C_{U,\boldsymbol{\mu}}$ and Φ_{μ} on $\boldsymbol{\mu}$ and instead use $\mathbb{H}$, C_U and Φ. Similarly, we will refer to $C_{U,\mu_k K_k}$ as $C_{U,k}$ and $\Phi_{\mu_k K_k}$ as Φ_k. We will also use the shorthand Π_U (resp. Π_S) instead of $P_r(C_U)$ (resp. $P_r(C_S)$).

Once a projection is defined, the labeled sample S is used to learn a linear hypothesis $x \mapsto \langle w, \Pi_U \Phi(x) \rangle_{\mathbb{H}}$ with bounded norm, $\|w\|_{\mathbb{H}} \leq 1$, in the subspace of the projected features $\Pi_U \Phi(x)$.

The two steps just described, dimensionality reduction by projection Π_U and supervised learning of w, may be coupled so that the best choice of weights $\boldsymbol{\mu}$ is made for the subsequent learning of w. (see Figure 3).

To do so, given constants $\Lambda_{(r)}$ and ν, we select vector $\boldsymbol{\mu}$ out of a convex set $\mathcal{M} =$

$$\left\{ \boldsymbol{\mu} \colon \|\boldsymbol{\mu}\|_{(r)} \leq \Lambda_{(r)}, \;\; \|\boldsymbol{\mu}\|_1 \leq 1, \;\; \sum_{k=1}^p \frac{1}{\mu_k} \leq \nu, \;\; \boldsymbol{\mu} \geq 0 \right\}.$$

We will show that the choice of this convex regularization set is crucial in guaranteeing the generalization ability our hypothesis class. The vector $\boldsymbol{\mu}$ is upper bounded by an L_1-norm inequality $\|\boldsymbol{\mu}\|_1 \leq 1$ (standard from learning kernels literature) but also by an inequality $\|\boldsymbol{\mu}\|_{(r)} \leq \Lambda_{(r)}$, where $\|\cdot\|_{(r)}$ is the semi-norm defined as the Ky Fan r-norm of C_U [Bhatia, 1997]:

$$\|\boldsymbol{\mu}\|_{(r)} = \|C_U\|_{(r)} = \sum_{i=1}^r \lambda_i(C_U). \tag{1}$$

The use of the semi-norm $\|\cdot\|_{(r)}$ in this context is key since $\|\boldsymbol{\mu}\|_{(r)}$ appears as the relevant quantity both in our generalization bounds and in our lower bounds. The lower bound constraint on $\boldsymbol{\mu}$, $\sum_{k=1}^{p} \mu_k^{-1} \leq \nu$, will imply an upper bound on the eigengap of the induced covariance operator, which is a fundamental quantity that influences the concentration of eigenspaces. In fact in Section 4 we give a simple example that demonstrates the dependence on the eigengap is tight, and also implies the necessity of the lower bound regularization.

Thus, the hypothesis set H defined by our supervised dimensionality reduction set-up is defined as follows:

$$H = \left\{ x \mapsto \langle w, \Pi_U \Phi(x) \rangle_{\mathbb{H}} : \|w\|_{\mathbb{H}} \leq 1, \boldsymbol{\mu} \in \mathcal{M} \right\}. \quad (2)$$

In the analysis that follows, we will also make use of normalized kernel matrices which, given a sample S of size m and kernel function K, are defined as $[\overline{\mathbf{K}}]_{i,j} = \frac{1}{m} K(x_i, x_j)$. Since the kernel matrix is normalized we have that $\lambda_i(\overline{\mathbf{K}}_k) = \lambda_i(C_{S,k})$, where $C_{S,k}$ is the sample covariance operator associated to K_k (see [Rosasco et al., 2010] Proposition 9.2). We also define the unscaled sample kernel matrix $[\mathbf{K}]_{i,j} = K(x_i, x_j)$.

We will assume that C_U admits at least r non-zero eigenvalues and will similarly assume that the set of kernel matrices $\overline{\mathbf{K}}_k$ of K_k associated to the sample U or S for any $k \in [1, p]$ contains at least one matrix with rank at least r. Furthermore, we assume that the base kernels K_k, $k \in [1, p]$, satisfy the condition $K_k(x, x) \leq 1$ for all $x \in \mathcal{X}$, which is guaranteed to hold for all normalized kernels. Finally, we assume the base kernels are *linearly independent with respect to the union of the samples S and U*.

Definition 1. Linearly Independent Kernels Let $K_1, \ldots, K_p$ be p PDS kernels and let $S = (x_1, \ldots, x_n)$ be a sample of size m. For any $k \in [1, p]$, let $\mathbb{H}_k$ denote the RKHS associated to K_k and $\overline{\mathbb{H}}_k$ the subspace of $\mathbb{H}_k$ spanned by the set of functions $\{\Phi_{K_k}(x_i) : i = 1, \ldots, m\}$. Then, $K_1, \ldots, K_p$ are said to be *linearly independent with respect to the sample S* if, for any $k \in [1, p]$, no non-zero function in $\overline{\mathbb{H}}_k$ can be expressed as a linear combination of the functions in $\cup_{l \neq k} \overline{\mathbb{H}}_l$.

Linear independence typically holds in practice, e.g., for polynomial and Gaussian kernels on $\mathbb{R}^N$. As an example, let $\mathcal{X} = \mathbb{R}^N$ and define the sample $S = \{x_1, \ldots, x_m\}$. Define two base kernels: Gaussian $K_1(x, y) = e^{-\|x-y\|^2}$ and linear $K_2(x, y) = \langle x, y \rangle$. Then $\Phi_{K_1}(x) : t \mapsto e^{-\|x-t\|^2}$, i.e. $\Phi_{K_1}(x)$ is an exponential function $e^{-\|x-t\|^2}$ with parameter x and argument t. In the same manner $\Phi_{K_2}(x) : t \mapsto \langle x, t \rangle$. Thus, $\overline{\mathbb{H}}_1$ is the span of exponential functions $\{e^{-\|x_1-t\|^2}, \ldots, e^{-\|x_m-t\|^2}\}$ and $\overline{\mathbb{H}}_2$ is the span of linear functions $\{\langle x_1, t \rangle, \ldots, \langle x_m, t \rangle\}$. Clearly, no exponential function can be represented as a linear combination of linear functions and likewise, in general, no linear function is represented as a (finite) linear combination of exponential functions. Thus, the base kernels K_1 and K_2 are linearly independent with respect to sample S as in Definition 2. Therefore $\overline{\mathbb{H}}_1 \perp \overline{\mathbb{H}}_2$ in the reproducing space of $K_1 + K_2$ and defined $\overline{\mathbb{H}} = \overline{\mathbb{H}}_1 \bigoplus \overline{\mathbb{H}}_2$. Such decomposition of $\overline{\mathbb{H}}$ into a direct sum allows to characterize the eigenfunctions of C_S: they consist of the eigenfunctions of $C_S|_{\overline{\mathbb{H}}_1}$ and $C_S|_{\overline{\mathbb{H}}_2}$. The eigenfunctions of each restricted operator are orthogonal to each other. Note that the orthogonality of eigenfunctions does not necessarily imply the orthogonality of the eigenvectors of sample kernel matrices $\overline{\mathbf{K}}_1$ and $\overline{\mathbf{K}}_2$.

More generally, the support of the base kernels can be straightforwardly modified to ensure that this condition is satisfied.

It follows from construction $\mathbb{H} = \mathbb{H}_1 + \cdots + \mathbb{H}_p$ and the results of [Aronszajn, 1950, Section 6] that when base kernels are linearly independent with respect to sample S, then $\overline{\mathbb{H}}_k$ are orthogonal subspaces of $\mathbb{H}$, thus we can define $\overline{\mathbb{H}} = \bigoplus_{k=1}^{p} \overline{\mathbb{H}}_k$, which will be extremely useful in decomposing the spectra of operators C_S. Linearly independent base kernels imply that C_S has at most pm nonzero eigenvalues of the form $\mu_k \lambda_j(C_{S,k})$, which is an explicit representation of the eigenvalues of C_S in terms of $\boldsymbol{\mu}$.

3 Generalization bound

In this section we outline the main steps taken in deriving a generalization bound as well as analyze the bound and discuss its implications. Proofs that are not included in this section can be found in the appendix.

The main result of this section is to derive an upper bound on the sample Rademacher complexity $\widehat{\mathfrak{R}}_S(H)$ of hypothesis class H. The sample Rademacher complexity of H is defined as $\widehat{\mathfrak{R}}_S(H) = \frac{1}{m} \mathrm{E}_{\boldsymbol{\sigma}} \left[\sup_{h \in H} \sum_{n=1}^{m} \sigma_n h(x_n) \right]$, where σ_n are i.i.d. random variables taking values +1 and -1 with equal probabilities. Once that is done we can then directly invoke the result of [Koltchinskii and Panchenko, 2002] and [Bartlett and Mendelson, 2003], which states that with probability at least $1 - \delta$ over the draw of sample $|S| = m$ and for all $h \in H$ the generalization error $R(h)$ is bounded by

$$R(h) \leq \widehat{R}_{S,\rho(h)} + \frac{2}{\rho} \widehat{\mathfrak{R}}_S(H) + 3 \sqrt{\frac{\log(2/\delta)}{2m}}, \quad (3)$$

where, given $c(x)$ is the true label of $x \in \mathcal{X}$, $R(h) = Pr_{x \sim \mathcal{D}}[h(x) \neq c(x)]$ and $\hat{R}_{S,\rho}(h)$ is the fraction of points in S with classification margin less than ρ.

Note, in our setting H is parametrized by w and $\boldsymbol{\mu}$ and that we can consider the supremum over these two parameters

separately. Finding the supremum over w can be done in a standard manner, using Cauchy-Schwarz,

$$\sup_{\|w\|\leq 1} \sum_{n=1}^{m} \sigma_n h(x_n) = \sup_{\|w\|\leq 1} \langle w, \Pi_U \sum_{n=1}^{m} \sigma_n \Phi(x_n) \rangle = \|\Pi_U \sum_{n=1}^{m} \sigma_n \Phi(x_n)\| .$$

Now it remains to compute $\sup_{\boldsymbol{\mu}\in\mathcal{M}} \|\Pi_U \sum_{n=1}^{m} \sigma_n \Phi(x_n)\|$, which is the more challenging expression. First of all, it will be more convenient to work with a projection defined with respect to C_S instead of C_U, since we are projecting instances from sample S. Similarly, we will find it useful to control a norm $\|C_S\|_{(r)}$ instead of $\|C_U\|_{(r)}$. Both of these issues can be addressed by using concentration inequalities to bound the difference of the projections Π_U and Π_S [Zwald and Blanchard, 2005] as well as the difference of the operators C_U and C_S [Shawe-Taylor and Cristianini, 2003]. For that we extend a constraint set $\mathcal{M}$ to a larger set $\mathcal{N}$:

$$\mathcal{N} = \Big\{ \boldsymbol{\mu} \colon \|C_S\|_{(r)} \leq \Lambda_{(r)} + \kappa, \ \ \|\boldsymbol{\mu}\|_1 \leq 1, \ \sum_{k=1}^{p} \frac{1}{\mu_k} \leq \nu, \ \ \boldsymbol{\mu} \geq 0 \Big\}, \quad (4)$$

where $\kappa = 4\Big(1+\sqrt{\frac{\log(2p/\delta)}{2}}\Big)$. In the following lemma we bound the transition from $\sup_{\boldsymbol{\mu}\in\mathcal{M}} \|\Pi_U \sum_{n=1}^{m} \sigma_n \Phi(x_n)\|$ to $\sup_{\boldsymbol{\mu}\in\mathcal{N}} \|\Pi_S \sum_{n=1}^{m} \sigma_n \Phi(x_n)\|$.

Lemma 3.1 *Let $C_{S,k}$ be the sample covariance operator of kernel $\mu_k K_k$ with a reproducing space $\mathbb{H}_k$. Define C_S (resp. C_U) as $C_S = \sum_{k=1}^{p} C_{S,k}$ and Π_S (resp. Π_U) be the orthogonal projection onto the eigenspace of $\lambda_i(C_S)$ for $i \in [1, r]$. Then with probability at least $1-\delta$ for any $u \in \mathbb{H} = \mathbb{H}_1 + \cdots + \mathbb{H}_p$*

$$\sup_{\boldsymbol{\mu}\in\mathcal{M}} \|\Pi_U u\| \leq \sup_{\boldsymbol{\mu}\in\mathcal{N}} \Big(\|\Pi_S u\| + \frac{8\kappa\nu}{\Delta_r \sqrt{m}} \|u\| \Big), \quad (5)$$

where $\Delta_r = \min_{k\in[1,p]} \big(\lambda_r(C_k) - \lambda_{r+1}(C_k)\big)$, C_k is the true covariance operator of kernel K_k and $\kappa = 4\Big(1 + \sqrt{\frac{\log(2p/\delta)}{2}}\Big)$.

Now using Lemma 3.1 and letting $u = \sum_{n=1}^{m} \sigma_n \Phi(x_n)$ we find that, with high probability, the Rademacher complexity of H is bounded by

$$\widehat{\mathfrak{R}}_S(H) \leq \frac{1}{m} \mathop{\mathrm{E}}_{\boldsymbol{\sigma}} \Bigg[\sup_{\boldsymbol{\mu}\in\mathcal{N}} \Big(\|\Pi_S \sum_{n=1}^{m} \sigma_n \Phi(x_n)\| + \frac{8\kappa\nu}{\Delta_r\sqrt{m}} \|\sum_{n=1}^{m} \sigma_n \Phi(x_n)\| \Big) \Bigg]. \quad (6)$$

We will distribute the supremum and bound each of the two terms separately. In the case of the second term, $\|\sum_{n=1}^{m} \sigma_n \Phi(x_n)\|$, we will upper bound the supremum over $\boldsymbol{\mu} \in \mathcal{N}$ with the supremum over a larger set constrained only by $\|\boldsymbol{\mu}\|_1 \leq 1$. This leads us to the expression $\frac{1}{m} \mathrm{E}_{\boldsymbol{\sigma}} \big[\sup_{\|\boldsymbol{\mu}\|_1\leq 1} \|\sum_{n=1}^{m} \sigma_n \Phi(x_n)\| \big]$, which is exactly equal to the Rademacher complexity of learning kernels for classification without projection. This complexity term can be bounded using Theorem 2 of [Cortes et al., 2010], which gives the following:

$$\frac{1}{m} \mathop{\mathrm{E}}_{\boldsymbol{\sigma}} \Bigg[\sup_{\|\boldsymbol{\mu}\|_1\leq 1} \|\sum_{n=1}^{m} \sigma_n \Phi(x_n)\| \Bigg] \leq \frac{\sqrt{\eta_0 e \lceil \log p \rceil}}{\sqrt{m}}, \quad (7)$$

where $\eta_0 = \frac{23}{22}$. Now it remains to bound the expectation of $\|\Pi_S \sum_{n=1}^{m} \sigma_n \Phi(x_n)\|$.

Lemma 3.2 *Let C_k and $C_{S,k}$ be the true and sample covariance operator of kernel K_k. Define C_S (resp. C_U) as $C_S = \sum_{k=1}^{p} C_{S,k}$ and let Π_S (resp. Π_U) be the orthogonal projection onto the eigenspace of $\lambda_i(C_S)$ (resp. $\lambda_i(C_{S'})$) for $i \in [1, r]$. Then with probability at least $1-\delta$.*

$$\frac{1}{m} \mathop{\mathrm{E}}_{\boldsymbol{\sigma}} \Bigg[\sup_{\boldsymbol{\mu}\in\mathcal{N}} \Big(\|\Pi_S \sum_{n=1}^{m} \sigma_n \Phi(x_n)\| \Big) \Bigg] \leq \frac{1}{\sqrt{m}} \sqrt{2\big(\Lambda_{(r)} + \kappa\big) \log(2p\,m)}, \quad (8)$$

where $\kappa = 4\Big(1 + \sqrt{\frac{\log(2p/\delta)}{2}}\Big)$.

Proof The term $\|\Pi_S \sum_{n=1}^{m} \sigma_n \Phi(x_n)\|$ is naturally bound using the constraint on $\|C_S\|_{(r)}$ since it involves projection onto eigenspace of C_S and $\|C_S\|_{(r)}$ controls its spectrum. Therefore we will reduce the problem to the supremum over $\|C_S\|_{(r)} \leq \epsilon$, where $\epsilon = \Lambda_{(r)} + \kappa$.

By Lemma B.1 (see appendix) we have that $\|\Pi_S \sum_{n=1}^{m} \sigma_n \Phi(x_n)\|^2 = m \mathbf{u}_\mu \cdot \mathbf{u}_\sigma$, where $\mathbf{u}_\mu$ is a vector with entries $\mu_k \lambda_j(\overline{\mathbf{K}}_k)$ and $\mathbf{u}_\sigma$ is a vector with entries $(\mathbf{v}_{k,j}^\top \boldsymbol{\sigma})^2$ such that $(k,j) \in I_\mu$. An indexing set I_μ is a set of pairs (k,j) that correspond to largest r eigenvalues of $\{\mu_k \lambda_j(\overline{\mathbf{K}}_k)\}_{k,j}$. In order remove the dependence of the indexing set on the identity of the top eigenvalues, we upper bound the expression over the choice of all size-r sets:

$$\sup_{\|C_S\|_{(r)}\leq\epsilon} \mathbf{u}_\mu \cdot \mathbf{u}_\sigma = \sup_{\|\mathbf{u}_\mu\|_1\leq\epsilon} \mathbf{u}_\mu \cdot \mathbf{u}_\sigma \leq \sup_{|I|=r} \Big(\sup_{\|\mathbf{u}_\mu\|_1\leq\epsilon} \mathbf{u}_\mu \cdot \mathbf{u}_\sigma \Big), \quad (9)$$

where $\sup_{|I|=r}$ indicates the supremum over all indexing sets I of size r. Then, by the dual norm property we have

$$\sup_{|I|=r} \sup_{\|\mathbf{u}_\mu\|_1\leq\epsilon} \mathbf{u}_\mu \cdot \mathbf{u}_\sigma = \sup_{|I|=r} \epsilon \|\mathbf{u}_\sigma\|_\infty = \epsilon \max_{k,j} (\mathbf{v}_{k,j}^\top \boldsymbol{\sigma})^2 . \quad (10)$$

Thus, $\|\Pi_S \sum_{n=1}^{m} \sigma_n \Phi(x_n)\|$ is bounded by the following:

$$\max_{k,j}\sqrt{m\epsilon(\mathbf{v}_{k,j}^\top\boldsymbol{\sigma})^2} \le \sqrt{m\epsilon}\max_{k,j}|\mathbf{v}_{k,j}^\top\boldsymbol{\sigma}| = \sqrt{m\epsilon}\max_{k,j}\max_{s_t\in\{-1,1\}} s_t\mathbf{v}_{k,j}^\top\boldsymbol{\sigma}\,.$$

By Massart's lemma [Massart, 2000]

$$\mathop{\mathrm{E}}_{\boldsymbol{\sigma}}\Big[\max_{k,j}\max_{s_t\in\{-1,1\}} s_t\mathbf{v}_{k,j}^\top\boldsymbol{\sigma}\Big] \le \sqrt{2\log(2pm)}\,. \tag{11}$$

This follows since the norm of $s_t\mathbf{v}_{k,j}$ is bounded by 1 and the cardinality of the set which the maximum is taken over is bounded by $2\sum_{k=1}^{p}\sum_{j=1}^{\mathrm{rank}(\overline{\mathbf{K}}_k)} \le 2pm$.

Combining all intermediate results brings us to the bound

$$\frac{1}{m}\mathop{\mathrm{E}}_{\boldsymbol{\sigma}}\left[\sup_{\|C_S\|_{(r)}\le\epsilon}\|\Pi_S\sum_{n=1}^{m}\sigma_n\Phi(x_n)\|\right] \le \frac{1}{\sqrt{m}}\sqrt{2\epsilon\log(2p\,m)}, \tag{12}$$

and the final result is obtained by letting $\epsilon = \Lambda_{(r)} + \kappa$. ∎

Thus, after combining Lemma 3.2 and 3.1 above we derive an upper bound on the expectation in (6), which gives us a bound on the sample Rademacher complexity. That bound is presented in the following theorem.

Theorem 3.3 *Let hypothesis set H be defined as in* (2)*. Then for any sample S of size $m < u$ drawn i.i.d. according to some distribution $\mathcal{D}$ over $\mathcal{X}\times\{-1,+1\}$ such that $\sqrt{m} > \frac{2\kappa}{\Delta_r}$ the empirical Rademacher complexity of the hypothesis set H can be bounded as follows with probability at least $1-\delta$,*

$$\widehat{\mathfrak{R}}_S(H) \le \frac{1}{\sqrt{m}}\left(\sqrt{2\big(\Lambda_{(r)}+\kappa\big)\log(2pm)} + \frac{8\kappa\nu\sqrt{\eta_0 e\lceil\log p\rceil}}{\Delta_r}\right), \tag{13}$$

where $\Delta_r = \min_k\big(\lambda_r(C_k) - \lambda_{r+1}(C_k)\big)$, $\kappa = 4\Big(1+\sqrt{\frac{\log(2p/\delta)}{2}}\Big)$ and $\eta_0 = \frac{23}{22}$.

If we consider only parameters p and $\Lambda_{(r)}$, then the Rademacher complexity bound in (13) is $O\Big(\sqrt{\frac{\Lambda_{(r)}\log(pm)}{m}}\Big)$. The learning scenario and regularization in standard learning kernels [Cortes et al., 2010] differs from ours, thus we will make a few adjustments that will allow us to compare those two bounds in a most coherent way, particularly, we will let $S = U$ and express $\Lambda_{(r)}$ in terms of unscaled sample kernel matrices.

$$\Lambda_{(r)} = \frac{1}{m}\sup_{|I|=r}\sum_{(k,j)\in I}\mu_k\lambda_j(\mathbf{K}_k) \le \frac{1}{m}\sup_{|I|=r}\sum_{(k,j)\in I}\lambda_j(\mathbf{K}_k) \tag{14}$$

That results in Rademacher complexity of $O\Big(\frac{1}{m}\sqrt{\sup_{|I|=r}\sum_{(k,j)\in I}\lambda_j(\mathbf{K}_k)\log(pm)}\Big)$, while the standard learning kernels bound is $O\Big(\frac{1}{m}\sqrt{\sup_{k\in[1,p]}\sum_{j=1}^{m}\lambda_j(\mathbf{K}_k)\log(p)}\Big)$. Here, $\sup_{|I|=r}\sum_{(k,j)\in I}\lambda_j(\mathbf{K}_k)$ is the largest r-long sum of eigenvalues that can be picked from any base kernel matrix, while $\sup_{k\in[1,p]}\sum_{j=1}^{m}\lambda_j(\mathbf{K}_k)$ is the largest $m-$long sum of eigenvalues that can be picked only from one base kernel matrix. Thus, if r is sufficiently smaller than m and given a certain choice of base kernels, learning kernels in the supervised dimensionality reduction problem will enjoy a tighter Rademacher complexity than standard learning kernels.

Finally, plugging in the upper bound from Theorem 3.3 into (3) results in the generalization bound of H. Note that the confidence term in (3) changes from $\log(2p/\delta)$ to $\log(4p/\delta)$, because Rademacher complexity is bounded with high probability. To the best of our knowledge, this is the first generalization guarantee provided for the use of projection in the reproducing space with a learned kernel in a supervised learning setting.

Theorem 3.4 *Let hypothesis set H be defined as in* (2)*. Then with probability at least $1-\delta$ over the draw of sample $|S| = m$ and for all $h \in H$ the generalization error $R(h)$ is bounded by*

$$\widehat{R}_{S,\rho}(h) + \frac{2}{\rho\sqrt{m}}\left(\sqrt{2\big(\Lambda_{(r)}+\kappa\big)\log(2pm)} + \frac{8\kappa\nu\sqrt{\eta_0 e\lceil\log p\rceil}}{\Delta_r}\right) + 3\sqrt{\frac{\log(4p/\delta)}{2m}}\,,$$

where $\hat{R}_{S,\rho}(h)$ is the fraction of points in S with classification margin less than ρ.

We note that [Mosci et al., 2007] and [Gottlieb et al., 2013] provide generalization bounds for supervised dimensionality reduction, however their learning scenarios are different from ours, particularly in a sense that they do not learn a mapping and projection for dimensionality reduction jointly with a discrimination function on reduced data. Nevertheless, their generalization bounds are comparable to ours in the special case $p = 1$. The analysis of [Gottlieb et al., 2013] is done for general metric spaces, while in the particular example of Euclidean space they show that generalization bound is $O\big(\sqrt{\frac{\kappa}{m}} + \sqrt{\frac{\eta}{m}}\big)$, where κ is the dimension of underlying data manifold and η is the average distance of training set to that manifold. Thus, while our bound has the same rate with respect to m as that of [Gottlieb et al., 2013], it is based on Ky-Fan semi-norm regularization $\Lambda_{(r)}$, while the bound of [Gottlieb et al., 2013] is based on the assumption that data approximately follows some low dimensional manifold with dimension

κ. The generalization bound in [Mosci et al., 2007] is $O\big(\frac{1}{\sqrt{m}}\big)$, however while we fix the number of eigenvalues for dimensionality reduction at r, their bound requires selecting all eigenvalues above a threshold $\lambda_m = O\big(\frac{1}{\sqrt{m}}\big)$. By applying this threshold they have control over the eigenspectrum of covariance operator, while we control the spectrum by the Ky-Fan regularization.

4 Lower bounds

In this section we show a lower bound on the Rademacher complexity of the hypothesis class H defined in (2). This lower bound demonstrates the central role that the Ky Fan r-norm regularization, $\|\mu\|_{(r)} \leq \Lambda_{(r)}$, plays in controlling the complexity of the hypothesis class. Furthermore, it demonstrates the tightness of the upper bound presented in the previous section in terms the number of training samples m. We additionally give a small example that demonstrates the necessity of the eigengap term which appears in Lemma 3.1 and which motivates the additional regularization term $\sum_{k=1}^{p} \mu_k^{-1} \leq \nu$ that is used to bound the eigengap in the proof of Lemma 3.1.

Theorem 4.1 *For any choice of m, r there exists samples S and U, a setting of the regularization parameter $\Lambda_{(r)}$, as well as a choice of base kernels $K_1, \ldots, K_p$ that guarantees*

$$\hat{\mathfrak{R}}_S(H) \geq \sqrt{\frac{\Lambda_{(r)}}{2m}}\,.$$

Proof First we let S and U be any two samples, both of size m, such that the U is simply an unlabeled version of S. Now, assume $p = 1$ and the sample kernel matrix $\overline{\mathbf{K}}_1$ of kernel K_1 has exactly r distinct non-zero simple eigenvalues. Finally, select $\Lambda_{(r)}$ such that $\Lambda_{(r)}/\lambda_1(\overline{\mathbf{K}}_1) \leq 1$.

As calculated in Section 3, $\sup_{\|w\|\leq 1} \sum_{n=1}^{m} \sigma_n h(x_n) = \|\Pi_U \sum_{n=1}^{m} \sigma_n \Phi(x_n)\|$ and in this particular scenario $\|C_U\|_{(r)} = \|C_S\|_{(r)}$, thus the empirical Rademacher complexity simplifies to $\hat{\mathfrak{R}}_S(H) = \frac{1}{m}\,\mathrm{E}_{\boldsymbol{\sigma}}\big[\sup_{\|C_S\|_{(r)}\leq\Lambda_{(r)}} \|\Pi_S \sum_{n=1}^{m} \sigma_n \Phi(x_n)\|\big]$, where the projection can be written directly in terms of the sample S and the L_1 constraint on $\boldsymbol{\mu}$ is not needed since it is satisfied by the Ky Fan r-norm constraint when $\Lambda_{(r)} \leq \lambda_1(\overline{\mathbf{K}}_1)$.

Now, let $u_1, \cdots, u_r$ denote the top r eigenfunctions of C_S, then following the steps from Lemma B.1 we can express the norm of projection as

$$\|\Pi_S \sum_{n=1}^{m} \sigma_n \Phi(x_n)\| = \sqrt{m \sum_{i=1}^{r} \lambda_i (\boldsymbol{\sigma}^\top \mathbf{v}_i)^2} \tag{15}$$

$$= \sqrt{m \sum_{j=1}^{r} \mu_1 \lambda_j(\overline{\mathbf{K}}_1)(\boldsymbol{\sigma}^\top \mathbf{v}_{1,j})^2}\,, \tag{16}$$

where $(\mu_1 \lambda_j(\overline{\mathbf{K}}_1), \mathbf{v}_{1,j})$ is the eigenpair of normalized sample kernel matrix $\mu_1 \overline{\mathbf{K}}_1 = \overline{\mathbf{K}}$. The expression is furthermore simplified by introducing the vectors $\mathbf{u}_\mu$ with entries $\mu_1 \lambda_j(\overline{\mathbf{K}}_1)$ and $\mathbf{u}_\sigma$ with entries $(\mathbf{v}_{1,j}^\top \boldsymbol{\sigma})^2$. Note that here, unlike in the general statement of Lemma B.1, the choice of r entries that appear in $\mathbf{u}_\mu$ and $\mathbf{u}_\sigma$ are not effected by the value of $\boldsymbol{\mu}$, since there are in fact only r nonzero eigenvalues total by construction (i.e. there is one base kernel of rank r). The choice of $\boldsymbol{\mu}$, however, still affects the scale of the r eigenvalues.

By the monotonicity of the square-root function and using the definition of $\mathbf{u}_{\boldsymbol{\mu}}$ as well as the dual norm we have

$$\sup_{\|C_S\|_{(r)}\leq\Lambda_{(r)}} \sqrt{\mathbf{u}_\mu \cdot \mathbf{u}_{\boldsymbol{\sigma}}} = \sqrt{\sup_{\|\mathbf{u}_{\boldsymbol{\mu}}\|_1 \leq \Lambda_{(r)}} \mathbf{u}_\mu \cdot \mathbf{u}_{\boldsymbol{\sigma}}} \tag{17}$$

$$= \sqrt{\Lambda_{(r)} \|\mathbf{u}_\sigma\|_\infty}\,. \tag{18}$$

Thus, the Rademacher complexity is reduced to

$$\begin{aligned} \hat{\mathfrak{R}}_S(H) &= \sqrt{\frac{\Lambda_{(r)}}{m}}\, \mathop{\mathrm{E}}_{\boldsymbol{\sigma}}\Big[\sqrt{\max_{j\in[1,r]} (\mathbf{v}_{1,j}^\top \boldsymbol{\sigma})^2}\Big] \\ &= \sqrt{\frac{\Lambda_{(r)}}{m}}\, \mathop{\mathrm{E}}_{\boldsymbol{\sigma}}\Big[\max_{j\in[1,r]} |\mathbf{v}_{1,j}^\top \boldsymbol{\sigma}|\Big]\,. \end{aligned} \tag{19}$$

Finally, we use Jensen's inequality and Khintchine's inequality to show

$$\mathop{\mathrm{E}}_{\boldsymbol{\sigma}}\Big[\max_{j\in[1,r]} |\mathbf{v}_{1,j}^\top \boldsymbol{\sigma}|\Big] \geq \max_{j\in[1,r]} \mathop{\mathrm{E}}_{\boldsymbol{\sigma}}\Big[|\mathbf{v}_{1,j}^\top \boldsymbol{\sigma}|\Big] \tag{20}$$

$$\geq \max_{j\in[1,r]} 2^{-1/2} \|\mathbf{v}_{1,j}\| = 2^{-1/2}, \tag{21}$$

where the tight constant $2^{-1/2}$ used in Khintchine's inequality can be found in Chapter II of [Nazarov and Podkorytov, 2000]. Plugging this constant back into equation (19) completes the theorem. ∎

The lower bound demonstrates the effect of the regularization parameter $\Lambda_{(r)}$ as well as the tightness of the upper bound in terms of m.

While Theorem 4.1 has shown the necessity of the Ky Fan r-norm constraint, we will now give a small example that illustrates the difference of projections (for example as seen in Lemma 3.1) must necessary depend on the

eigengap quantity. This in turn motivates the regularization $\sum_{k=1}^{p} \frac{1}{\mu_k} \leq \nu$ which ensures that eigengap is not arbitrarily small, since otherwise if $\boldsymbol{\mu}$ goes to zero, then Δ_r also goes to zero. The fact that the eigengap is essential for the concentration of projections has been known in the matrix perturbation theory literature [Stewart and Sun, 1990]. The following proposition gives an example which shows that the dependence on the eigengap is tight.

Proposition 4.2 *There exists operators A and B such that*

$$\|P_r(A) - P_r(B)\| = \frac{2\|A-B\|}{\lambda_r(A) - \lambda_{r+1}(A)} .$$

where $P_r(A)$ (resp. $P_r(B)$) is the orthogonal projection onto the top r eigenspace of A (resp. B).

Proof Consider $r = 1$ and A and B defined as follows $A = \left(\begin{smallmatrix} 1+\epsilon & 0 \\ 0 & 1 \end{smallmatrix}\right)$ and $B = \left(\begin{smallmatrix} 1 & 0 \\ 0 & 1+\epsilon \end{smallmatrix}\right)$, thus $A - B = \left(\begin{smallmatrix} \epsilon & 0 \\ 0 & -\epsilon \end{smallmatrix}\right)$, which implies that $\|A - B\| = \epsilon$. Also, the eigengap is equal to $\lambda_1(A) - \lambda_2(A) = \epsilon$. Now note that $P^1(A)$ is the projection onto $e_1 = (1, 0)^\top$ and $P^1(B)$ is the projection onto $e_2 = (0, 1)^\top$. Since e_1 and e_2 are orthogonal, this implies $\|P^1(A) - P^1(B)\| = \|P^1(A)\| + \|P^1(B)\| = 2$. On the other hand, $\frac{2\|A-B\|}{\lambda_1(A)-\lambda_2(A)} = \frac{2\epsilon}{\epsilon} = 2$, which complete the proof. ∎

Assume that operator C is defined together with a positive $\mu \in \mathbb{R}^d$, then the stability of the r-eigenspace of μC is determined by $\frac{1}{\mu} \frac{1}{\lambda_r(C)-\lambda_{r+1}(C)}$. When we have operators $C_1, \ldots, C_p$, with identical spectra such that $\lambda_r(C_k) - \lambda_{r+1}(C_k) = \epsilon$ for each k in $[1, p]$, but each of them acting on mutually orthogonal subspaces, the stability of the r-eigenspace of $C = \sum_{k=1}^{p} \mu_k C_k$ is determined by $\sum_{k=1}^{p} \frac{1}{\mu_k} \frac{1}{\lambda_r(C_k)-\lambda_{r+1}(C_k)} = \frac{1}{\epsilon} \sum_{k=1}^{p} \frac{1}{\mu_k}$.

This example clearly shows that $\sum_{k=1}^{p} \frac{1}{\mu_k}$ controls the stability of eigenspace when operators act on orthogonal subspaces, which directly applies to linearly independent kernels.

5 Discussion

Here, we briefly discuss the results presented. Let us first emphasize that our choice of the hypothesis class H (Section 2) is strongly justified a posteriori by the learning guarantees we presented: both our upper and lower bounds on the Rademacher complexity (Sections 3 and 4) suggest controlling the quantities present in the definition of H. The regularization parameters we provide can be tuned to directly bound each of these crucial quantities and thereby limit the risk of over fitting.

Second, we observe that the hypothesis class H clearly motivates the design of a single-stage coupled algorithm. Such an algorithm would be based on structural risk minimization (SRM) and seek to minimize the empirical error over increasingly complex hypothesis sets, by varying the parameters $\Lambda_{(r)}$ and ν, to trade-off empirical error and model complexity. Although the design and evaluation of such an algorithm is beyond the scope of this paper, we note that existing literature has empirically evaluated both learning kernels with KPCA in an unsupervised (two-stage) fashion [Zhuang et al., 2011, Lin et al., 2011] and applied dimensionality reduction (single-stage training) with a fixed kernel function [Fukumizu et al., 2004], [Gönen, 2014]. While these existing algorithms do not directly consider the hypothesis class we motivated, they can, in certain cases, still select a hypothesis function that is found in our class. In particular, our learning guarantees remain applicable to hypotheses chosen in a two-stage manner, as long as the regularization constraints are satisfied. Similarly the case $p = 1$ which corresponds to the standard fixed-kernel supervised learning scenario is covered by our analysis. We note that even in such cases, the bounds that we provide would be the first to guarantee the generalization ability of the algorithm via bounding the sample Rademacher complexity.

6 Algorithm

In this section we obtain a computational expression for $h(x)$, where $x \in \mathbb{R}^d$ and $h \in H$. Moreover, we formulate a minimization problem for training $\boldsymbol{\mu}$ and w as well as discuss ways to efficiently solve it by breaking into a series of convex sub-problems.

For the clarity of presentation we assume Π_U is full rank and provide the expression for $h \in \mathbb{H}$ in the following lemma.

Lemma 6.1 *Let $x \in \mathbb{R}^d$, then for every $h \in H$ there exist real numbers $\{z_{k,j}\}_{k\in[1,p],j\in[1,m]}$ such that*

$$h(x) = \sum_{k,j} \xi_{k,j}(\boldsymbol{\mu}) z_{k,j} \sqrt{\frac{\mu_k}{\lambda_j(\mathbf{K}_k)}} \sum_{n=1}^{m} K_k(x, x_n)[\mathbf{v}_{k,j}]_n \tag{22}$$

where

$$\xi_{k,j}(\boldsymbol{\mu}) = \begin{cases} 1 \text{ if } \mu_k\lambda_j(\mathbf{K}_k) \text{ in top } r \text{ from } \{\mu_k\lambda_j(\mathbf{K}_k)\}_{k,j} \\ 0 \text{ otherwise} \end{cases} \tag{23}$$

with the following constraints

$$\sum_{k,j} z_{k,j}^2 \leq 1 \tag{24}$$

$$\frac{1}{m} \sum_{k,j} \xi_{k,j}(\boldsymbol{\mu}) \mu_k \lambda_j(\mathbf{K}_k) \leq \Lambda_{(r)} \tag{25}$$

$$\sum_k \mu_k \leq 1 \tag{26}$$

$$\sum_k \frac{1}{\mu_k} \leq \nu \tag{27}$$

$$\mu_k \geq 0 \tag{28}$$

Proof First, to obtain the computational expression for $h(x_n) = \langle w, \Pi_U \Phi(x) \rangle_{\mathbb{H}}$, for the moment imagine that Π_U is of full rank, then if $z_1(w), \cdots, z_{mp}(w)$ are the coordinates of w in the span of eigenfunctions of C_U and $z_1(\Pi_U \Phi(x)), \cdots, z_{mp}(\Pi_U \Phi(x))$ are the coordinates of $\Pi_U \Phi(x)$ in that span, we will have $\langle w, \Pi_U \Phi(x) \rangle_{\mathbb{H}} = \sum_{j=0}^{mp} z_j(w) z_j(\Pi_U \Phi(x))$. Now, when we go back to the original scenario of rank r projection Π_U, we introduce choice variables $\xi_1, \cdots, \xi_{mp} \in 0, 1$, where $\xi_i = 1$ if the i-th eigenfunction is chosen for projection and $\xi_i = 0$ otherwise. Thus, the expression for $h(x)$ becomes

$$h(x) = \sum_{j=0}^{mp} \xi_j z_j(w) z_j(\Pi_U \Phi(x)) \tag{29}$$

The assumption of linearly independent kernels allows us to break the sum above into p components for each base kernel. We will do it by keeping two indices $k \in [1, p]$ and $j \in [1, m]$, which gives $h(x) = \sum_{k,j} \xi_{k,j} z_{k,j}(w) z_{k,j}(\Pi_U \Phi(x))$. Observe that $z_{k,j}(\Pi_U \Phi(x)) = \langle \Phi(x), u_{k,j} \rangle = u_{k,j}(x)$ and the steps in the proof of Lemma B.1 show that

$$u_{k,j}(x) = \sqrt{\frac{\mu_k}{\lambda_j(\mathbf{K}_k)}} \sum_{n=1}^{m} K_k(x, x_n) [\mathbf{v}_{k,j}]_n \tag{30}$$

Moreover, varying w in $\mathbb{H}$ for the purpose of our algorithm is equivalent to varying its coordinates $z_{k,j}(w)$, thus we will use variables $z_{k,j}$ instead of them with the constraint $\sum_{k,j} z_{k,j}^2 \leq 1$. Given all this analysis, the computational expression for $h(x)$ becomes

The optimization variables are $z_{k,j}$ and μ_k, while $\xi_{k,j}$ is determined by $\boldsymbol{\mu}$. ∎

When we define $\mathbf{z}$ as vector with entries $z_{k,j}$ and have some convex loss function over the training sample $L(\boldsymbol{\mu}, \mathbf{z})$, the optimization problem is

$$\min_{\boldsymbol{\mu}, \mathbf{z}} L(\boldsymbol{\mu}, \mathbf{z}) \tag{31}$$

subject to

$$\|\mathbf{z}\| \leq 1 \tag{32}$$

$$\boldsymbol{\mu} \in \mathcal{M} \tag{33}$$

Note that the loss function includes the complicated term

$$\sum_{k,j} \xi_{k,j}(\boldsymbol{\mu}) z_{k,j} \sqrt{\frac{\mu_k}{\lambda_j(\mathbf{K}_k)}} \sum_{n=1}^{m} K_k(x, x_n) [\mathbf{v}_{k,j}]_n \,. \tag{34}$$

For conciseness, define

$$c_{k,j}(x) = \sqrt{\frac{1}{\lambda_j(\mathbf{K}_k)}} \sum_{n=1}^{m} K_k(x, x_n) [\mathbf{v}_{k,j}]_n \,, \tag{35}$$

which allows us to write clearly $h(x, \mathbf{z}, \boldsymbol{\mu}) = \sum_{k,j} \xi_{k,j} c_{k,j}(x) z_{k,j} \sqrt{\mu_k}$. We will make a substitution $w_{k,j} = z_{k,j} \sqrt{\mu_k}$. That substitution changes constraint $\|\mathbf{z}\| \leq 1$ to $\sum_{k,j} \frac{w_{k,j}^2}{\mu_k} \leq 1$, which is a convex set in w and $\boldsymbol{\mu}$ for $\mu_k > 0$.

This reduces the problem to the following constrained optimization problem

$$\min \frac{1}{n} \sum_{i=1}^{n} L\left(\sum_{k,j} \xi_{k,j} c_{k,j}(x_n) w_{k,j}, y_n \right) \tag{36}$$

subject to

$$\sum_{k,j} \frac{w_{k,j}^2}{\mu_k} \leq 1 \tag{37}$$

and

$$\boldsymbol{\mu} \in \mathcal{M} \tag{38}$$

There are at least two possible ways to relax the problem. First, we can relax $\xi_{k,j}(\boldsymbol{\mu})$ to no longer be a function of $\boldsymbol{\mu}$ and make $\xi_{k,j} \in \{0, 1\}$ a discrete optimization variable instead with an additional constraint $\sum_{k,j} \xi_{k,j} = r$. The second option is even weaker: we can let $\xi_{k,j}$ be a continuous variable in the interval $[0, 1]$ and keep the constraint $\sum_{k,j} \xi_{k,j} = r$. Investigation of those relaxation steps is a direction for further research.

7 Conclusion

We presented a new analysis and generalization guarantees for the scenario of coupled nonlinear dimensionality reduction with a learner kernel. The hypothesis class is designed with regularization constraints that are directly motivated by the generalization guarantee, which we show lower bounds for as well. Our analysis invites the design of learning algorithms for selecting hypotheses from this specifically tailored class, either in a two-stage or a single-stage manner.

References

Anthony, M. and Bartlett, P. (1999). *Neural network learning: theoretical foundations*. Cambridge University Press.

Aronszajn, N. (1950). Theory of reproducing kernels. *Transactions of the American mathematical society*, pages 337–404.

Bartlett, P. L. and Mendelson, S. (2003). Rademacher and gaussian complexities: Risk bounds and structural results. *The Journal of Machine Learning Research*, 3:463–482.

Belkin, M. and Niyogi, P. (2001). Laplacian eigenmaps and spectral techniques for embedding and clustering. In *NIPS*, volume 14, pages 585–591.

Bhatia, R. (1997). *Matrix Analysis*. Springer.

Blanchard, G. and Zwald, L. (2008). Finite-dimensional projection for classification and statistical learning. *Information Theory, IEEE Transactions on*, 54(9):4169–4182.

Cortes, C., Mohri, M., and Rostamizadeh, A. (2009). L 2 regularization for learning kernels. In *Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence*, pages 109–116.

Cortes, C., Mohri, M., and Rostamizadeh, A. (2010). Generalization bounds for learning kernels. In *Proceedings of the 27th International Conference on Machine Learning (ICML-10)*, pages 247–254.

Dhillon, P. S., Foster, D. P., Kakade, S. M., and Ungar, L. H. (2013). A risk comparison of ordinary least squares vs ridge regression. *The Journal of Machine Learning Research*, 14(1):1505–1511.

Fukumizu, K., Bach, F. R., and Jordan, M. I. (2004). Dimensionality reduction for supervised learning with reproducing kernel hilbert spaces. *The Journal of Machine Learning Research*, 5:73–99.

Gönen, M. (2014). Coupled dimensionality reduction and classification for supervised and semi-supervised multilabel learning. *Pattern recognition letters*, 38:132–141.

Gönen, M. and Alpaydın, E. (2011). Multiple kernel learning algorithms. *The Journal of Machine Learning Research*, 12:2211–2268.

Gottlieb, L.-A., Kontorovich, A., and Krauthgamer, R. (2013). Adaptive metric dimensionality reduction. In *Algorithmic Learning Theory*, pages 279–293. Springer.

Ham, J., Lee, D. D., Mika, S., and Schölkopf, B. (2004). A kernel view of the dimensionality reduction of manifolds. In *Proceedings of the twenty-first international conference on Machine learning*, page 47. ACM.

Hegde, C., Wakin, M., and Baraniuk, R. (2008). Random projections for manifold learning. In *Advances in neural information processing systems*, pages 641–648.

Kloft, M., Brefeld, U., Sonnenburg, S., and Zien, A. (2011). Lp-norm multiple kernel learning. *The Journal of Machine Learning Research*, 12:953–997.

Koltchinskii, V. and Panchenko, D. (2002). Empirical margin distributions and bounding the generalization error of combined classifiers. *Annals of Statistics*, pages 1–50.

Lanckriet, G. R., Cristianini, N., Bartlett, P., Ghaoui, L. E., and Jordan, M. I. (2004). Learning the kernel matrix with semidefinite programming. *The Journal of Machine Learning Research*, 5:27–72.

Lin, Y.-Y., Liu, T.-L., and Fuh, C.-S. (2011). Multiple kernel learning for dimensionality reduction. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 33(6):1147–1160.

Massart, P. (2000). Some applications of concentration inequalities to statistics. *Annales de la Faculté des Sciences de Toulouse*, 9(2):245–303.

Mohri, M., Rostamizadeh, A., and Talwalkar, A. (2012). *Foundations of machine learning*. MIT press.

Mosci, S., Rosasco, L., and Verri, A. (2007). Dimensionality reduction and generalization. In *Proceedings of the 24th international conference on Machine learning*, pages 657–664. ACM.

Nazarov, F. L. and Podkorytov, A. N. (2000). Ball, haagerup, and distribution functions. In *Complex analysis, operators, and related topics*, pages 247–267. Springer.

Pearson, K. (1901). Liii. on lines and planes of closest fit to systems of points in space. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 2(11):559–572.

Rosasco, L., Belkin, M., and Vito, E. D. (2010). On learning with integral operators. *The Journal of Machine Learning Research*, 11:905–934.

Roweis, S. T. and Saul, L. K. (2000). Nonlinear dimensionality reduction by locally linear embedding. *Science*, 290(5500):2323–2326.

Shawe-Taylor, J. and Cristianini, N. (2003). Estimating the moments of a random vector with applications. In *Proceedings of the GRETSI 2003 Conference*, pages 47–52.

Stewart, G. and Sun, J. (1990). *Matrix Perturbation Theory*. Academic Press.

Tenenbaum, J. B., De Silva, V., and Langford, J. C. (2000). A global geometric framework for nonlinear dimensionality reduction. *Science*, 290(5500):2319–2323.

Zhuang, J., Wang, J., HOI, C. H., and Lan, X. (2011). Unsupervised multiple kernel learning. *Journal of Machine Learning Research (JMLR)*, 20:129–144.

Zwald, L. and Blanchard, G. (2005). On the convergence of eigenspaces in kernel principal component analysis. In *Proceedings of NIPS*, pages 1649–1656.

A Proof of Lemma 3.1

Proof For the first part of the proof, let C_{μ_k} be the true covariance operator of kernel $\mu_k K_k$. Since for each k both $C_{S,k}$ and $C_{U,k}$ approach C_{μ_k} with high probability, we will show a concentration bound on their difference that holds uniformly over $k \in [1,p]$ as well as U and S. Using union bound for probabilities and Lemma 1 from [Zwald and Blanchard, 2005] (equivalently Corollary 5 from [Shawe-Taylor and Cristianini, 2003]) with probability at least $1-\delta$ for all $k \in [1,p]$,

$$\max\left[\|C_{\mu_k} - C_{S,k}\|_{\overline{\mathbb{H}}_k}, \|C_{\mu_k} - C_{U,k}\|_{\overline{\mathbb{H}}_k}\right] \leq 2\mu_k M_\delta/\sqrt{m}\,, \tag{39}$$

where $M_\delta = 1 + \sqrt{\frac{\log(2p/\delta)}{2}}$. We used $u > m$ to obtain the bound.

By triangle inequality and decomposition over orthogonal subspaces of $\overline{\mathbb{H}} = \bigoplus_{k=1}^p \overline{\mathbb{H}}_k$, the norm $\|\Pi_S - \Pi_U\|$ is bounded by $\sum_{k=1}^p \|P_r(C_{\mu_k}) - P_r(C_{S,k})\|_{\overline{\mathbb{H}}_k} + \sum_{k=1}^p \|P_r(C_{\mu_k}) - P_r(C_{U,k})\|_{\overline{\mathbb{H}}_k}$. By Theorem 3[2] from [Zwald and Blanchard, 2005] for each $k \in [1,p]$,

$$\|P_r(C_{\mu_k}) - P_r(C_{S,k})\|_{\overline{\mathbb{H}}_k} \leq \frac{8\|C_{\mu_k} - C_{S,k}\|_{\overline{\mathbb{H}}_k}}{\lambda_r(C_{\mu_k}) - \lambda_{r+1}(C_{\mu_k})}\,, \tag{40}$$

and a similar statement holds for projeciton with respect to sample U.

We will use $\|\boldsymbol{\mu}\|_1 \leq 1$ to make the bound in (39) be $M_\delta/\sqrt{m}$ and we will decompose $\lambda_r(C_{\mu_k}) - \lambda_{r+1}(C_{\mu_k}) = \mu_k\big(\lambda_r(C_k) - \lambda_{r+1}(C_k)\big) \geq \mu_k \Delta_r$, where C_k is the true covariance operator of kernel K_k and $\Delta_r = \min_{k\in[1,p]}\big(\lambda_r(C_k) - \lambda_{r+1}(C_k)\big)$. Now $2M_\delta/\sqrt{m}\mu_k\Delta_r$ is the uniform bound on the norm of projections in (40). Summing up $\|P_r(C_{\mu_k}) - P_r(C_{S,k})\|_{\overline{\mathbb{H}}_k} + \|P_r(C_{\mu_k}) - P_r(C_{U,k})\|_{\overline{\mathbb{H}}_k}$ over k and applying the uniform bound $2M_\delta/\sqrt{m}\mu_k\Delta_r$, which holds for both samples U and S, we conclude

$$\|\Pi_S - \Pi_U\| \leq \sum_{k=1}^p \frac{32M_\delta}{\mu_k\Delta_r\sqrt{m}} \leq \frac{32M_\delta\nu}{\Delta_r\sqrt{m}} \tag{41}$$

For the last part we use a simple series of inequalities to get

$$\big|\|C_U\|_{(r)} - \|C_S\|_{(r)}\big| \leq \sum_{i=1}^r |\lambda_i(C_U) - \lambda_i(C_S)| \tag{42}$$

$$\leq \sqrt{r}\left(\sum_{i=1}^r |\lambda_i(C_U) - \lambda_i(C_S)|^2\right)^{1/2}, \tag{43}$$

which is in turn bounded by $\sqrt{r}\|C_U - C_S\|$ using Hoffman-Wielandt inequality. Now, $\|C_U - C_S\|$ is simply bounded by $\sum_{k=1}^p \|C_{\mu_k} - C_{S,k}\|_{\overline{\mathbb{H}}_k} + \sum_{k=1}^p \|C_{\mu_k} - C_{U,k}\|_{\overline{\mathbb{H}}_k}$. If we apply the uniform bound from (39) in the form $\mu_k M_\delta/\sqrt{m}$, we get that with probability at least $1-\delta$

$$\big|\|C_U\|_{(r)} - \|C_S\|_{(r)}\big| \leq \sum_{k=1}^p \frac{\sqrt{r}\mu_k 4M_\delta}{\sqrt{m}} \leq 4M_\delta\,. \tag{44}$$

Putting together the two main bounds, we have that:

$$\sup_{\boldsymbol{\mu}\in\mathcal{M}} \|\Pi_U u\| \leq \sup_{\boldsymbol{\mu}\in\mathcal{N}} \|\Pi_U u\| \leq \sup_{\boldsymbol{\mu}\in\mathcal{N}} \left(\|\Pi_S u\| + \frac{32M_\delta\nu}{\Delta_r\sqrt{m}}\|u\|\right), \tag{45}$$

∎

B Lemma B.1 with proof

Lemma B.1 *Let $C_{S,k}$ be the sample covariance operator of kernel $\mu_k K_k$. Define C_S (resp. C_U) as $C_S = \sum_{k=1}^p C_{S,k}$ and Π_S (resp. Π_U) be the orthogonal projection onto the eigenspace of $\lambda_i(C_S)$ for $i \in [1,r]$. Let $\mathbf{v}_{k,j}$ be the eigenvector corresponding to $\lambda_j(\overline{\mathbf{K}}_k)$. If I_μ is an indexing set that contains pairs (k,j) that correspond to largest r eigenvalues from the set $\{\mu_k\lambda_j(\overline{\mathbf{K}}_k)\}_{k,j}$, then*

$$\|\Pi_S \sum_{n=1}^m \sigma_n \Phi(x_n)\|^2 = m \sum_{(k,j)\in I_{\boldsymbol{\mu}}} \mu_k\lambda_j(\overline{\mathbf{K}}_k)(\mathbf{v}_{k,j}^\top \boldsymbol{\sigma})^2 \tag{46}$$

$$= m\mathbf{u}_\mu \cdot \mathbf{u}_\sigma\,, \tag{47}$$

where $\mathbf{u}_\mu$ is a vector with entries $\mu_k\lambda_j(\overline{\mathbf{K}}_k)$, and $\mathbf{u}_\sigma$ is a vector with entries $(\mathbf{v}_{k,j}^\top\boldsymbol{\sigma})^2$ such that $(k,j) \in I_\mu$.

Proof By properties of orthogonal projection, $\|\Pi_S \sum_{n=1}^m \sigma_n\Phi(x_n)\|^2$ can be expressed as $\sum_{i=1}^r \|P_{\lambda_i}(C_S)\sum_{n=1}^m \sigma_n\Phi(x_n)\|^2$, where $P_{\lambda_i}(C_S)$ is the orthogonal projection onto the eigenfunction that corresponds to $\lambda_i(C_S)$. Recall from Section 2 that when base kernels are linearly independent with respect to sample S, then $\overline{\mathbb{H}} = \bigoplus_{k=1}^p \overline{\mathbb{H}}_k$, which implies that for every $i \in [1,m]$ there exists a (k,j) such that the eigenfunction of $\lambda_i(C_S)$ is equal to the eigenfunction of $\mu_k\lambda_j(C_{S,k})$. Note that j may not be equal to i since $\boldsymbol{\mu}$ influences the ordering of eigenvalues. Thus, we define the indexing set $I_{\boldsymbol{\mu}}$ that contains the r pairs of indices (k,j) that correspond to the r largest eigenvalues in $\{\mu_k\lambda_j(C_{S,k})\}_{k,j}$ for particular value of $\boldsymbol{\mu}$.

For any $f \in \overline{\mathbb{H}}$ we can express $\|P_{\lambda_j}(C_S)f\|^2 = \langle u_j, f\rangle^2$, where u_j is the eigenfunction corresponding to $\lambda_j(C_S)$. Since $\overline{\mathbb{H}} = \bigoplus_{k=1}^p \overline{\mathbb{H}}_k$ we can express the norm of projection as $\|\Pi_S f\|^2 = \sum_{(k,j)\in I_{\boldsymbol{\mu}}}\langle u_{k,j}, f\rangle^2$, where $u_{k,j}$ is an eigenfunction of $C_{S,k}$ with eigenvalue $\lambda_j(C_{S,k})$. When

[2] Note that the actual theorem we reference works provided that $2\|C_{\mu_k} - C_{S,k}\|_{\overline{\mathbb{H}}_k}/\big(\lambda_r(C_{\mu_k}) - \lambda_{r+1}(C_{\mu_k})\big) \leq 1/2$, which stems from the convergence of power series used in the proof. Here, for the simplicity of presentation we multiply their bound by 4, which relaxes such a requirement

$f = \sum_{n=1}^{m} \sigma_n \Phi(x_n)$ observe that $u_{k,j}$ belongs to orthogonal component $\overline{\mathbb{H}}_k$, therefore it suffices to take inner product in $\overline{\mathbb{H}}_k$, which by the reproducing property is equal to $u_{k,j}(x_n)$. By [Blanchard and Zwald, 2008, Equation 18] $u_{k,j}(x_n)$ takes the form

$$u = \sqrt{\frac{\mu_k}{\lambda_j(\overline{\mathbf{K}}_k)m}} \sum_{i=1}^{m} K_k(\cdot, x_i)[\mathbf{v}_{k,j}]_i \,, \tag{48}$$

which results in the following series of equalities

$$\begin{aligned}
&\langle u_{k,j}, \sum_{n=1}^{m} \sigma_n \Phi(x_n) \rangle = \\
&= \sum_{n=1}^{m} \sigma_n \langle u_{k,j}, \Phi(x_n) \rangle \\
&= \sqrt{m} \sqrt{\frac{\mu_k}{\lambda_j(\overline{\mathbf{K}}_k)}} \sum_{n=1}^{m} \sigma_n \sum_{i=1}^{m} \left[\mathbf{v}_{k,j}\right]_i \frac{1}{m} K_k(x_n, x_i) \\
&= \sqrt{m} \sqrt{\frac{\mu_k}{\lambda_j(\overline{\mathbf{K}}_k)}} \lambda_j(\overline{\mathbf{K}}_k) \sum_{n=1}^{m} \sigma_n \left[\mathbf{v}_{k,j}\right]_n \\
&= \sqrt{m} \sqrt{\mu_k \lambda_j(\overline{\mathbf{K}}_k)} \mathbf{v}_{k,j}^{\top} \boldsymbol{\sigma} \,.
\end{aligned}$$

Squaring the terms above and summing them up, we arrive at $\|\Pi_S \sum_{n=1}^{m} \sigma_n \Phi(x_n)\|^2 = m \sum_{(k,j)\in I_\mu} \mu_k \lambda_j(\overline{\mathbf{K}}_k)(\mathbf{v}_{k,j}^{\top}\boldsymbol{\sigma})^2 = m\mathbf{u}_\mu \cdot \mathbf{u}_\sigma$ which complete the proof. ∎